Introduction to Supernetting

Supernetting allows the use of multiple IP networks on the same interface. It is the reverse of subnetting, which allows the use of a single IP network on multiple interfaces. Another document in the REDNET Knowledgebase describes subnetting.

Officially, supernetting is the term used when multiple network addresses of the same Class are combined into blocks. If the IP networks are contiguous, you may be able to use a supernet. If the IP networks are not contiguous, you would need to use sub-interfaces. These are not currently supported on Compatible Systems routers but are supported on routers from Cisco Systems.

A prerequisite for supernetting is that the network addresses be consecutive and that they fall on the correct boundaries. To combine two Class C networks, the first address' third octet must be evenly divisible by 2. If you would like to supernet 8 networks, the mask would be 255.255.248.0 and the first address' third octet needs to be evenly divisible by 8. For example, 198.41.15.0 and 198.41.16.0 could NOT be combined into a supernet, but you would be able to combine 198.41.18.0 and 198.41.19.0 into a supernet.

An IP address is a 32-bit number (4 bytes, called "octets", separated by periods, commonly called "dots.") Supernetting is most often used to combine Class C addresses (the first octet has values from 192 through 223). A single Class C IP network has 24 bits for the network portion of the IP address, and 8 bits for the host portion of the IP address. This gives a possibility of 256 hosts within a Class C IP network (2^8=256).

The subnet mask for a Class C IP network is normally 255.255.255.0. To use a supernet, the number of bits used for the subnet mask is REDUCED. For example, by using a 23 bit mask (255.255.254.0 -- 23 bits for the network portion of the IP network, and 9 bits for the host portion), you effectively create a single IP network with 512 addresses. Supernetting, or combining blocks of IP networks, is the basis for most routing protocols currently used on the Internet.

For Example: Two Class "C" network numbers of 198.41.78.0 and 198.41.79.0

The addresses pass the prerequisites. They are consecutive and the third octet of the first address is divisible by 2 (78 Mod 2 = 0). To further illustrate what is being done, let's look at the addresses in binary. The third octet of the first address (78) is 01001110. The second (79) is 01001111. The binaries are the same except for the last bit of the address (the 24th bit of the IP address). The 78 network is supernet 0 and the 79 network is supernet 1.

The subnet mask for this example supernet is 23 bits, or 255.255.254.0. ALL devices on the network MUST be using this subnet mask. Any device that is not using this subnet mask would be unreachable.

The broadcast address for ALL devices on the example supernet is 198.41.79.255. Most modern devices don't require you to fill out the broadcast address, as it can be deduced from the IP address and the subnet mask. The broadcast address is used as a special destination signifying ALL hosts on the network.

As with any IP network, the first number in the range (.0 in a class "C") has special significance, and can't be assigned to any hosts on the network. The first number in the range is referred to as the "network number". Conversely, the last, or highest number in the range (.255 in a class "C") is called the broadcast address, and also can't be used by any host on the network.

Because of these unique addresses, it would probably be wise not to use the 198.41.78.255 and 198.41.79.0 addresses (in the above example), even though these SHOULD be perfectly legal